EXHIBIT I

www.archive.org
415.561.6767
415.840-0391 e-fax

mail:
Internet Archive
PO Box 29244
San Francisco, CA
94129-0244

ship:
Internet Archive
116 Sheridan Avenue
Presidio of San Francisco
San Francisco, CA 94129

AFFIDAVIT OF PAUL HICKMAN

1. I am the Office Manager at the Internet Archive, located at the Presidio of San Francisco, California. I make this declaration of my own personal knowledge.

2. The Internet Archive is a website that provides access to a digital library of Internet sites and other cultural artifacts in digital form. Like a paper library, we provide free access to researchers, historians, scholars, and the general public. The Internet Archive is affiliated with and receives support from various institutions, including the Library of Congress.

3. The Internet Archive has created a service known as the Wayback Machine. The Wayback Machine makes it possible to surf more than 55 billion pages stored in the Internet Archive's web archive. Visitors to the Wayback Machine can type in a URL (i.e., a website address), select a date range, and then begin surfing on an archived version of the Web. The links on the archived files, when served by the Wayback Machine, point to other archived files (whether HTML pages or images). If a visitor clicks on a link on an archived page, the Wayback Machine will serve the archived file with the closest available date to the originally requested page.

4. The Internet Archive receives data from third parties who compile the data by using software programs known as crawlers that surf the Web and automatically store copies of website files at certain points in time as they existed at that point in time. This data is donated to the Internet Archive, which preserves and provides access to it.

5. The Internet Archive assigns a URL on its site to the archived files in the format http://web.archive.org/web/[Year in yyyy][Month in mm][Day in dd][Time code in hh:mm:ss]/[Archived URL]. Thus, the Internet Archive URL http://web.archive.org/web/19970126045828/http://www.archive.org/ would be the URL for the record of the Internet Archive home page HTML file (http://www.archive.org/) archived on January 26, 1997 at 4:58 a.m. and 28 seconds (1997/01/26 at 04:58:28). Typically, a printout from a Web browser will show the URL in the footer. The date assigned by the Internet Archive applies to the HTML file but not to image files linked therein. Thus images that appear on the printed page may not have been archived on the same date as the HTML file. Likewise, if a website is designed with "frames," the date assigned by the Internet Archive applies to the frameset as a whole, and not the individual pages within each frame.

6. Attached hereto as Exhibit A are true and accurate copies of printouts of the Internet Archive's records of the HTML files archived from the URLs and the dates specified in the footer of the printout.

7. I declare under penalty of perjury that the foregoing is true and correct.

DATE: 4/17/06

Paul Forrest Hickman

Golden State Notary Acknowledgment Form

State of California
County of San Francisco. } ss.

On April #17th before me, David Ryan,
personally appeared Paul Forrest Hickman

personally known to me (or proved to me on the basis of satisfactory evidence) to be the person(s) whose name(s) is/are subscribed to the within instrument and acknowledged to me that he/she/they executed the same in his/her/their authorized capacity(ies), and that by his/her/their signature(s) on the instrument the person(s), or the entity upon behalf of which the person(s) acted, executed the instrument.

WITNESS my hand and official seal.

DAVID RYAN
COMM. # 1546323
NOTARY PUBLIC-CALIFORNIA
SAN FRANCISCO COUNTY
My Commission Expires
February 19, 2010

Signature of Notary

Affidavit Paul Hickman. Notes Acknowledgment.

Please provide information about the document that this form is attached to
This is not required under California State notary public law.

Last updated: May 30, 1997, 4:00 pm

TABLE OF CONTENTS

Q22: How can I configure RealSecure specifically for my network?

Q23: Can RealSecure be used for URL blocking?

Q24: Can I customize RealSecure's response to a network event?

Q25: Can RealSecure flag SSH and SSL traffic?

Q26: Can RealSecure log and flag the type and size of traffic or network service?

Q27: How does RealSecure detect a SYN flood?

Q28: How many RealSecure engines can a RealSecure console manage at one time?

Q29: Can RealSecure data be analyzed with a decision support system?

Q30: How are updates handled? Can an administrator upgrade fifty engines across an enterprise (for example) without losing configuration settings?

Q31: Can multiple RealSecure engines run on a single host with multiple adapter cards?

Q32: This product gathers a lot of information about my network. How should the RealSecure host be configured in order to protect this product from misuse?

Q33: How do I get a copy of RealSecure?

Q34: Whom do I contact for technical support?

Q35: Whom do I contact with product suggestions?

Q1: What is RealSecure?

A: RealSecure(TM) is a real-time, automated attack recognition and response system. It sits on your network, monitoring the network traffic stream looking for attacks and unauthorized access attempts. When RealSecure detects an attack, it can respond in a variety of ways, including logging the connection, notifying the network administrator, and killing the connection automatically.

Back To Top

Q2: What kinds of network events does RealSecure recognize?

A: RealSecure recognizes two types of network occurrences:

Attacks

Network activity patterns indicating that someone may be engaged in unauthorized or undesirable activity involving the systems and/or data on your network. Examples of these include SATAN scans, ping floods, WinNuke packets, SYN floods, IP half scans, and attempts to obtain unauthorized root access.

Sessions

Non-attack network activity that may be of interest to the network administrator. Examples of these include HTTP activity (who's surfing the net and where they're going), analysis of access to Windows shares (e.g., connections from engineering to accounting), and e-mail session decoding.

RealSecure is shipped with the most comprehensive set of attack recognition patterns in the industry.

Back To Top

Q3: What types of protocols can RealSecure see and decode?

A: RealSecure can filter and monitor any TCP/IP protocol. The network administrator can configure RealSecure to filter by protocol (TCP, UDP, ICMP), source port, destination port, source IP address, and/or destination IP address. RealSecure can interpret the following network services: web surfing, e-mail, file transfer, remote login, chat, talk and a host of others. The range of services that RealSecure can analyze is extended regularly, so be sure to check the ISS web site at http://www.iss.net for the latest status.

In addition, RealSecure will soon be able to monitor and decode Microsoft CIFS/SAMBA traffic for Windows networking environments.

Back To Top

Q4: How does RealSecure work?

A: RealSecure is installed on a host having a network adapter card. RealSecure puts the adapter card in promiscuous mode so that it acts like a sniffer and receives all the traffic on the local network segment. If a packet meets the filter criteria currently in force, it is parsed by the decode and attack recognition logic. Each active session is maintained and tracked, so that attack patterns that span many packets can be detected. This way, when an "interesting event" is detected, the appropriate actions can be taken.

RealSecure is completely unobtrusive. It only monitors the local traffic. RealSecure does not add any delay to the network segment.

Back To Top

Q5: How does RealSecure respond to attacks?

A: The actions taken upon detection of an attack or unauthorized activity are determined by the administrator. The administrator may choose from the following options:

- Display a message indicating that the event occurred
- View the session in real-time (or record for later playback)
- Kill the connection automatically by sending a reset packet to each session participant
- E-mail a notification to the administrator
- Execute a user-specified program
- Log the data related to the event for later reporting or playback

Back To Top

Q6: What platforms can RealSecure run on?

A: Currently, RealSecure is supported on the following platforms:

- SunOS 4.1.3 or later
- Linux 1.3.x
- Solaris (SPARC) 2.3 or later

However, RealSecure for Windows NT platforms is currently scheduled to ship in July, 1997.

Back To Top

Q7: What networks can RealSecure monitor?

A: RealSecure currently operates over Ethernet networks only. ISS will be adding support for FDDI and Fast Ethernet in the short term, with support for additional networking topologies to follow.

Back To Top

Q8: What Ethernet cards does RealSecure support?

A: RealSecure will operate over any Ethernet card that is capable of supporting promiscuous mode. Check the documentation for your Ethernet adapter to determine whether your card has this capability.

Back To Top

Q9: What are the recommended specifications required for a host to run RealSecure?

A: SunOS and Solaris:	SPARC 5 or better At least 32 Mbytes of RAM At least 60 Mbytes of free disk space Motif installation (Solaris 2.3, 2.4 for GUI) Ethernet interface connected to the target network
Linux:	Pentium 90 MHz. or better At least 32 Mbytes of RAM At least 60 Mbytes of free disk space X-Window system, version 11 or higher for GUI Ethernet interface connected to the target network
Windows NT:	Pentium 90 MHz. or better At least 32 Mbytes of RAM At least 60 Mbytes of free disk space Windows NT 4.0

Ethernet interface connected to the target network.

The configurations listed above assume that the management console and an engine are running on the specified host. If you are running an engine only, then you can reduce the RAM requirements to 16 Mbytes.

Back To Top

Q10: How is RealSecure deployed across the enterprise network?

A: RealSecure uses a distributed architecture. The RealSecure engine performs its filtering and monitoring functions on a given network segment. The RealSecure management console displays and logs the data and acts as a centralized engine management point.

Many RealSecure engines can report to a single management console. As engines detect unauthorized activity they take the appropriate action and then send a message to the management console so that the administrator can see what has happened. Engines can also upload their local log files and databases to the management console periodically, so that the network administrator has a centralized report of network activity.

With regard to placement of RealSecure engines, the best rule is to place a RealSecure engine on each segment where there is critical data to protect or a set of users that should be monitored.

Note that a RealSecure engine will only see the traffic that is on the local network segment. Since routers prevent traffic from being copied to inappropriate segments, several RealSecure engines might be needed for complete coverage of network activity.

The following figure shows a sample deployment of RealSecure.

- There is a RealSecure engine behind the firewall monitoring the traffic flow on and off the network. This engine will detect unauthorized activity from the Internet. It will also help the administrator detect misconfigurations in the firewall.
  Many companies also place a RealSecure engine outside the firewall in the DMZ, to protect the external web server and to analyze external traffic.
- There is a RealSecure engine on the Sales subnet because this is where the sales database resides.
- The RealSecure engine on the Marketing subnet protects the business plans and marketing strategies on the systems in that department.
- The Engineering subnet also supports a RealSecure engine because the source code archive is kept there.
- All of these engines can report to a single management console. That console might be located on the company's internal network or at a headquarters across the internet or at a Network Operations Center staffed by a service organization.

In this sample deployment, it is also possible to install RealSecure engines on the backbone that would see all the intersegment traffic. However, installing engines on each segment provides two distinct advantages:

- Each engine can be customized for the local segment. For example, the engine on the

sales subnetwork can be configured to examine all traffic directed to the sales database.
- Engines on the local segments will detect unauthorized activity that is initiated on that segment, an engineer attempting to gain root access to the source code archive, for example.

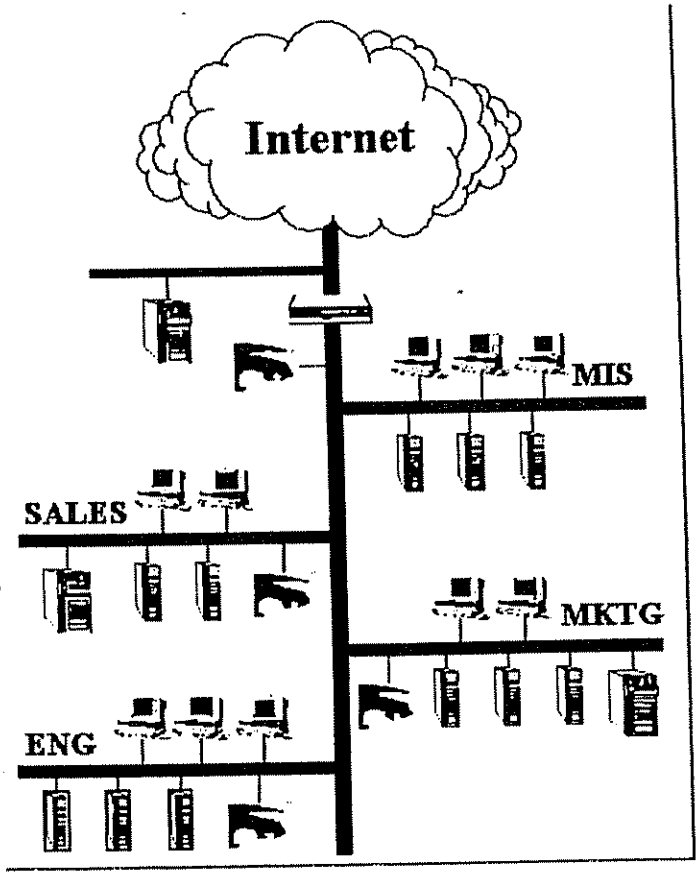

Back To Top

Q11: How do the RealSecure engines communicate with the RealSecure management console?

A: Data from the engines to the management console includes:

- Event messages - indications that something interesting has happened. These are passed up to the management console as they occur.
- Raw session data - the keystroke or data content of a session. This information is passed up to the management console as it occurs if the action associated with an event is "View Session".
- Database and log file information. These are sent up to the management console on demand.

Data from the management console to the engines includes:

- Start, stop, and pause commands.

- Changes to filter rules, attack signatures, and event responses.
- Keep-alive checks.
- Software updates.

The engines and the management console communicate using TCP port 590 and UDP ports 900 and 901. Data passed between the systems is encrypted and authenticated.

Back To Top

Q12: What authentication scheme do you use? What encryption scheme do you use?

A: The authentication scheme is a variation of the challenge-handshake authentication protocol. The management console provides a list of the engines that it manages along with a pass phrase used to authenticate data from each engine. Each engine contains a similar list for the management console. When one side wants to send a message to the other, the pass phrase is appended to the data, an MD5 checksum is calculated for the entire data set, and the checksum is attached to the packet (without the pass phrase). The entire message is encrypted and sent out over UDP.

Upon receiving the message, the other side of the connection will decrypt the data field, remove the checksum, attach the pass phrase, calculate an MD5 checksum, and compare with the checksum received. If they match, the message is authenticated.

The current encryption scheme is a fully exportable ISS proprietary encryption method. RealSecure will use a standard encryption method in a subsequent release.

Back To Top

Q13: How does RealSecure differ from a firewall? Don't they do the same things?

A: Firewalls and RealSecure use similar technologies to accomplish different things. Firewalls are *controlling* entities. They enforce general entry and exit rules for an entire network and aren't designed to look for attack patterns. Their main purpose is to keep the wrong kind of traffic off the network and their definition of "wrong kind of traffic" is usually based on IP address or protocol type.

RealSecure is not a product that controls network access. RealSecure does not interfere with the network traffic stream at all. Instead, it allows the traffic to go by and quietly watches for signs of unauthorized activity. RealSecure's definition of "unauthorized activity" is a sophisticated and customizable database of attack signatures.

Think of your network as a high-rise apartment building, the firewall as the doorman, and each RealSecure engine as a guard dog on a specific floor. The doorman is generally pretty good about letting the right people in and keeping the wrong people out. However, a moderately clever criminal will be able to get past the doorman and into the building. The guard dog is better at knowing who's authorized to be on the floor and responding quickly to stop the intrusion.

Back To Top

Q14: Do I need firewalls if I have RealSecure?

A: Absolutely. RealSecure is an essential addition to, but not a replacement for, your firewall security. When firewalls are properly configured, they keep out most undesired traffic. However, in order to provide some level of access, firewalls have tunnels and these tunnels can be exploited by would-be attackers. A good example is FTP. Many companies have an FTP server inside their network and associated tunnels through the firewall to allow access. A common attack is to attempt to gain root access to the FTP server. Once an attacker has access to a system inside the network, other systems become vulnerable. And although the firewall will not stop this type of attack, RealSecure will. By monitoring the traffic stream on the network behind the firewall, RealSecure can detect and terminate attempts to gain root access on the FTP server.

The other unfortunate reality is that firewalls are often misconfigured. A poorly configured firewall offers about as much protection as cheap sunglasses on an August afternoon. Although firewall misconfigurations should be fixed as soon as possible, having RealSecure inside the network can catch much of the undesirable traffic that's leaking through. Even if you choose not to terminate these undesired connections, the sheer number of alarms that RealSecure will generate will quickly indicate that your firewall is not doing its job.

Back To Top

Q15: What do I have to do to my network to run RealSecure?

A: Nothing. You don't have to buy new routers or bridges; you don't have to install any software on your servers; you don't have to reconfigure any devices. RealSecure works with your existing network infrastructure. All you need to do is place a UNIX® or Windows NT® system with an Ethernet card on the segment to be monitored and install RealSecure.

Back To Top

Q16: Will RealSecure run on a switched network?

A: Yes. Right now, this is primarily an issue of deployment.

Network switches (as well as IP switches) break a network into segments. Traffic that is not addressed to a system on a given segment will not be transmitted on that segment by a switch. Therefore, one alternative is to place an engine on each segment that contains critical data that need to be monitored carefully.

There are switches, however, that allow only a single MAC address to be attached to each port. These are usually referred to as "desktop switches". If you have desktop switches, it will be necessary to deploy a RealSecure engine at a higher level of the network, upstream from the switch, so that the engine will detect traffic between the switch and the rest of the network.

It's also important to note that ISS is working on ways to provide more detailed RealSecure coverage of switched networks.

Back To Top

Q17: How much delay does RealSecure add to the network?

A: None. Unlike firewalls, which often store and evaluate data before forwarding it to the inner network, RealSecure is completely unobtrusive. RealSecure monitors the network traffic, copying packets as needed, but does not alter or delay the traffic at all. The only time that RealSecure will have any impact the traffic flow is when it terminates connections in response to an attack and this won't be noticed, except by the attacker.

Back To Top

Q18: How much additional traffic does RealSecure add to the network?

A: In a standalone configuration (i.e., engine and management console on the same host), RealSecure will add no additional traffic to the network, since all communication between engine and management console takes place on the host.

In a distributed configuration (i.e., engines distributed around the enterprise network reporting data to the management console), the amount of additional traffic will depend on several factors:

- The number and frequency of network events reported from the engine to the management console.
- The frequency of database and log file uploads from the engine to the console.
- The size of the database and log file uploads from the engine to the management console. The administrator may choose not to upload everything, but may want a subset of the stored information instead.

Back To Top

Q19: Can RealSecure be completely transparent? Must RealSecure have an IP address?

A: Yes. RealSecure can be completely transparent. It is possible to monitor a network without the knowledge of the users on the network.

The RealSecure engine requires TCP/IP to communicate with the management console. Consequently, it requires an IP address. This is true even when the engine and management console are running on the same host.

Back To Top

Q20: Can RealSecure detect unauthorized activity in a Windows networking environment?

A: The next versions (Unix in June 1997, Windows NT in July 1997) will include the ability to decode SAMBA/ CIFS protocols for Windows networking. The product will also include several new attack signatures specific to the Windows networking environment. These will include the ability to detect:

- When one user attempts to copy a password file (.pwl) from a shared volume on a Windows 95 system.

- Remote registry access attempts.
- Null sessions.
- Attempts to read from or write to protected shares.

Back To Top

Q21: Can RealSecure play back logged traffic data at a later date?

A: RealSecure cannot playback logged data as if it were being received from the adapter card. However, network events can be stored in log files and databases for retrieval at a later date. RealSecure provides sophisticated reporting features that allow the administrator to sort and format event data by priority, source address, destination address, or network service over some period of time.

RealSecure also includes the ability to record the raw, binary content of an entire network session. This data is stored in a log file and can be replayed through the management console interface. It is played back exactly as it was received, keystroke for keystroke, so that the administrator can see how the attack or session unfolded.

Back To Top

Q22: How can I configure RealSecure specifically for my network?

A: There are two ways to customize RealSecure.

First, you can add your own filter rules to RealSecure. RealSecure can be instructed to filter on any combination of the following:

- Protocol
- Source IP address
- Destination IP address
- Source port
- Destination port

For example, a network administrator might want to log all traffic to the server that contains the financial data for the corporation. She would do this by adding a filter that would catch all traffic to the IP address of the server.

Second, you can alter the actions that RealSecure takes when an attack or event is detected. These actions include the following:

- Log a summary of the event (date, time, source, target, type of event).
- Log the entire binary content of a session.
- Display a message on the management console about the event.
- Notify the network administrator via e-mail.
- View the session in real-time.
- Kill the connection.
- Execute a user-defined program (UNIX version only).

Back To Top

Q23: Can RealSecure be used for URL blocking?

A: Yes, but in a limited fashion. RealSecure is not designed to be a "network nanny" that enforces appropriate usage policies on a network. Its real function is security.

You can install filters that match certain web sites and you can instruct RealSecure to terminate connections that match these filters. For example, I might install a new filter that terminates all connections from 208.21.4.0 (sexkitten.com) and from port 80 (HTTP). However, RealSecure is not a product that is designed to be used in this manner and you will find that there are other, easier ways of blocking certain URLs on your LAN.

Back To Top

Q24: Can I customize RealSecure's response to a network event?

A: Yes. See question 22.

Back To Top

Q25: Can RealSecure flag SSH and SSL traffic?

A: Yes.

Back To Top

Q26: Can RealSecure log and flag the type and size of traffic or network service?

A: RealSecure can log and flag the type of traffic, but not the size. RealSecure is not designed to monitor or manage the performance of the network segment, but the security.

Back To Top

Q27: How does RealSecure detect a SYN flood?

A: Some of the attacks that RealSecure detects involve more than just a single packet or single protocol type. Some involve variables that can be tuned for your network. SYN Flood is a good example. A SYN Flood is a denial of service attack. When a TCP connection is established, the initiator of the connection (the attacker, in this example) sends a SYN packet to the destination (the target system, in this example). The target system will acknowledge the connection and allocate memory to hold information about the connection. By establishing, but not using, many TCP connections, the attacker can cause the target machine to run out of memory and possibly crash.

RealSecure detects SYN Floods by monitoring the TCP connections that are established and by setting thresholds for the number of outstanding connections on a given machine at a given time. The network administrator may adjust the value of this threshold as appropriate for the network. There are several other attack signatures, like SYN Flood, that have tunable

parameters.

Back To Top

Q28: How many RealSecure engines can a RealSecure console manage at one time?

A: There is no hard and fast limitation to the number of engines that can be controlled by a single RealSecure management console. The practical number depends on several factors:

- The system configuration of the host running the management console software.
- The amount of traffic that flows between the engine and the console.
- The number of attacks and events recognized by the engine.

For example, a console managing twenty engines on a busy network with lots of attacks in progress can receive more data than a console managing 100 engines on a quiet network with few attacks and very tight filter rules.

In addition, there is also a human limit as to how many subnetworks can be managed from a single point. Practically speaking, the number of engines that will normally be report to a single console will depend on the geographic and organizational limitations of the controlling organization.

Back To Top

Q29: Can RealSecure data be analyzed with a decision support system?

A: Yes, if the decision support system is capable of reading an ODBC database.

Back To Top

Q30: How are updates handled? Can an administrator upgrade fifty engines across an enterprise (for example) without losing configuration settings?

A: Updates are posted on the ISS web site (http://www.iss.net) and users are notified of the new software via e-mail. The new release can be downloaded and installed at the administrator's convenience. Installation uses built-in file copy capabilities and is as simple as copying a new executable to the appropriate locations.

Since configuration settings (i.e., filter rules, enabled attacks, engines being managed) are saved in separate configuration files, installation of new software will have no effect on the current settings.

Back To Top

Q31: Can multiple RealSecure engines run on a single host with multiple adapter cards?

A: Not at present. There is currently a limit of one RealSecure engine per host. However, this is a feature that will be supported in a subsequent release.

Back To Top

Q32:This product gathers a lot of information about my network. How should the RealSecure host be configured in order to protect this product from misuse?

A: RealSecure is an amazingly powerful tool designed for network administrators. However, it could become a potentially dangerous tool in the wrong hands. It can grab user names, passwords, and even e-mail and file transfer content. Therefore, ISS recommends the following:

- Scan the engine and management console with ISS' Internet Scanner and System Security Scanner (S3) to minimize the system's vulnerability to attack. Use RealSecure on a dedicated host. Do not run any other applications on the system.
- Disable all services except for TCP/IP. The RealSecure engine reads raw data link packets from the adapter card, but uses TCP and UDP to communicate with the management console.
- Ensure that nothing is listening on any of the ports except for TCP 590 and UDP 900 and 901. These are the ports used for engine-console communication.
- Ensure that root or administrator access to the device is restricted. It would be a good idea if all other logins were disabled.

Back To Top

Q33: How do I get a copy of RealSecure?

A: Download an evaluation copy from the ISS web site at http://www.iss.net. Or, you can call us at 1-800-PROBE-62 (1-800-776-2362).

Back To Top

Q34: Whom do I contact for technical support?

A: You can send e-mail to support@iss.net. Or you can download our tech support FAQ from the ISS web site at http://www.iss.net. Finally, you can call us at 1-800-776-2362 and ask for technical support.

Back To Top

Q35: Whom do I contact with product suggestions?

A: For RealSecure, send an e-mail to the Product Manager, Mark Wood, at mwood@iss.net.

Back To Top

Technical Support: support@iss.net

Disclaimer:

The information contained in this FAQ may change without notice. Use of this information constitutes acceptance for product usage in an "AS IS" condition. There are NO warranties with regard to this information. In no event shall ISS be liable for any damages whatsoever arising out of, or in connection with the use or spread of this information. Any use of this information is at the user's own risk.

04/30/97

www.archive.org
415.561.6767
415.840-0391 e-fax

mail:
Internet Archive
PO Box 29244
San Francisco, CA
94129-0244

ship:
Internet Archive
116 Sheridan Avenue
Presidio of San Francisco
San Francisco, CA 94129

The following document is the .pdf found at
http://web.archive.org/web/19970709202307/www.iss.net/prod/rs_wp/realtime.pdf

It is 15 pages long.

Paul Forrest Hickman

Real-time attack recognition and response:

A solution for tightening network security

41 Perimeter Center
Suite 660
Atlanta, GA 30346
(800) 776-2362
www.iss.net

Table of Contents

Executive summary

Enterprise networks are becoming more difficult to secure for several important reasons. As organizations connect their local area networks (LANs) into wide area enterprise networks, these networks become more complex and, therefore, more difficult to secure. In an effort to share information and streamline operations, organizations are also opening their networks to business partners, suppliers, and other outsiders. These open networks are more susceptible to attack than their predecessors. In addition, organizations are connecting their internal networks to the Internet to reap the benefits of its assorted services and nearly universal reach. Connecting to the Internet exposes internal networks to millions of outsiders and greatly increases the difficulty of maintaining effective security.

Technology vendors are responding with a variety of security solutions to help organizations protect their internal networks from outside attacks. These solutions include firewalls, operating system security mechanisms such as authentication and access privilege levels, and encryption. Even with this combination of security solutions in place, hackers still manage to penetrate. Complicating the problem of maintaining effective security is the fact that networks are continually changing to meet shifting business situations, such as reorganizations, acquisitions, and mergers.

What is required is a security solution that is independent of conventional security mechanisms—one that detects and intercepts security breaches that penetrate the network's first line of defense. One such solution is an *attack recognition and response system*.

Typically implemented in software, an attack recognition and response system continually monitors network traffic, looking for known patterns of attack. When it detects an unauthorized activity, the software responds automatically with predetermined actions. It may report the attack, log the event, or terminate the unauthorized connection. Attack recognition and response software operates in concert with other security mechanisms to provide truly effective security.

This paper presents the concepts of network monitoring, attack recognition and response software. It describes the need for additional network security, and explains how attack recognition and response software addresses this need. It also examines the requirements for an effective attack recognition and response system. Finally, it describes RealSecure, a real-time attack recognition and response system that meets these requirements.

Networks are more vulnerable to attack

Organizations are continuing to expand their enterprise networks—connecting and consolidating local area networks into wide area networks (WANs). The resulting increase in complexity makes security more difficult and increases the vulnerability of the network to attack by both external and internal users. In addition, organizations are continuing to open their networks to outsiders. An example of this would be interconnecting enterprise networks with those of outside organizations. This further increases vulnerability to attack.

The Internet increases vulnerability

Today, many organizations are connecting their internal networks to the Internet to meet important business objectives that include:

- *Giving employees access to Internet resources.* Employees can increase their productivity by taking advantage of the vast amount of information and services on the Internet.
- *Giving external users access to the internal network from the Internet.* Organizations need to make internal network information available to users in the outside world, including customers, suppliers, and business partners.
- *Using the Internet as a basis for commerce.* One of the major attractions of the Internet is that it enables organizations to reach customers in much greater numbers and in far more locations than with conventional commerce vehicles.
- *Using the Internet as wide area network backbone.* The Internet provides an economical medium for connecting local area networks into WANs.

While connecting to the Internet offers numerous advantages, it also exposes internal networks to millions of outsiders. This exposure intensifies the need for comprehensive security.

Security must be enhanced

Increasing network complexity, greater openness, and the growing emphasis on the Internet are causing organizations to feel more and more insecure about their networks—and rightly so. These three trends are resulting in significantly higher exposure to both internal and external attacks. Ernst & Young/ Information Week's fourth annual Information Security Survey from 10/21/96 reveals some eye-opening statistics:

> "78% of Information Security chiefs, Information Security Officers and other high-level executives lost money from security breeches. More than 25% reported losses greater than twenty-five thousand dollars, and inside hackers were at fault for nearly 32%. With these breeches in security and serious financial losses, 70% have no more than three people dedicated to corporate security."

A variety of security devices are being deployed to protect internal networks from outside attacks. One type of security device receiving a considerable amount of attention today is the firewall. A firewall is typically implemented in software. It interposes a barrier at the point of connection between the Internet and the corporate internal network to keep out attackers. A recent survey in InfoWeek showed that 20 percent of respondents already have a firewall and over 40 percent more plan to implement one.

Firewalls, however, are not foolproof. They are difficult to configure, even by experts. They require the accurate configuration of numerous and confusing access control lists. Even a minor error in entering configuration information can result in a gaping hole in security. Complicating the issue is that firewall configurations must be continually updated to allow access to new network services required by end-users and to keep up with changing security policies. These policy changes are driven by changing business situations, such as reorganizations, mergers, and acquisitions. As a result of the complexity and continual change of firewall configurations, administrators are apt to miss potential security holes.

Networks provide an ideal medium to encourage collaboration among people. As a result, organizations are demanding network software that allows an increasing richness of collaboration and interaction among people. Of course, implied in this demand is that security not be compromised by increased levels of collaboration. As a result, configuring firewalls will become even more difficult as organizations increase their level of interaction with partners, suppliers, and customers.

Even properly configured firewalls have known weak spots. Using techniques such as IP spoofing and IP fragmentation, hackers have demonstrated the ability to pass through a large percentage of the firewalls on the market today. Another problem is that, in many situations, hackers can circumvent firewalls entirely. For example, if an internal user connects a modem to a networked PC and forgets to associate a password with the modem line, a hacker could enter the internal network directly through that modem, completely bypassing the firewall. Internal attacks from inside the firewall are also a frequent source of break-ins. In fact, internal attacks, often committed by disgruntled employees or co-opted contractors, account for the majority of network break-ins. FBI reports show that more than 60 percent of computer crimes originate inside the enterprise.

In addition to add-on devices such as firewalls, security mechanisms are also built into operating systems. Operating systems provide user authentication through passwords, and they provide multi-level access control to information. Like firewalls, however, operating system security is also vulnerable to attack. Operating systems are difficult to configure from a security perspective. Another problem is that software updates to operating systems can introduce security holes that are unknown to the administrator. Additionally, access control at the operating system level does not necessarily map directly to the network level. As a result, it isn't always possible for operating system access control to help block attacks from the network.

The rapid increase of network complexity and the attendant increase in the difficulty of securing networks has caused numerous vulnerabilities that far exceed an organization's capacity to deal with them effectively. To make matters worse, new vulnerabilities are being introduced into the networking environment every day. The result is a widening gap between an organization's security policy and its actual security practice. Organizations need solutions that help close this gap by augmenting traditional security systems with enhanced security mechanisms.

Closing the gap

There are two effective methods for augmenting conventional security systems to close the gap between security policy and security practice:

- *Attack recognition and response software.* This software continually monitors network traffic, looking for known patterns of attack. It runs on network machines that are strategically located at control points throughout the network , such as near an Internet router link or near a LAN on which critical data resides. When the software detects unauthorized activity, it responds automatically with some form of defined action. These actions can typically be configured by the administrator. Attack recognition and response software is like a burglar alarm. The alarm detects an intrusion in process and responds automatically by sounding an audible signal or telephoning the police.

- *Security scanning software.* This software probes the network with a series of tests to ferret out potential security vulnerabilities. It produces a detailed report on the weak spots that it finds, with sufficient information to enable corrective action. Security scanning software is similar to a security consultant. The consultant examines the organization's facilities looking for security weak spots and reports on the weaknesses discovered along with the necessary corrective actions to minimize or eliminate them.

This paper focuses on attack recognition and response software.

Attack detection

Attack recognition and response software typically detects attacks using one of the following two approaches:

- *Rule-based.* This approach draws from a library of known attack patterns or unauthorized activity and watches for those specific types of attacks. This is similar to the technique used in virus detection. The attack pattern library is updated continually as new types of attacks are discovered.

- *Statistical anomaly.* This approach operates on the assumption that users and networks always exhibit a predictable pattern of behavior and do not depart from this pattern over short periods of time. A deviation is considered to be an attack.

Attack response

Attack recognition and response software can be configured to react automatically to an attack in a variety of ways. It can:

- Log the event along with associated information.
- Alert appropriate personnel through console messages, e-mail, or pagers.
- Terminate the offending connection.
- Call a user-defined script or program.
- Perform a combination of these actions.

Combining software tools to enhance security

Attack recognition and response software can operate in conjunction with security scanning software to provide more complete protection. For example, a vulnerability may appear temporarily and then disappear between security scans. As a result, the scanning software does not detect it. The attack recognition and response software, however, does detect an attack through that vulnerability. The software reports the attack and logs information about it. The administrator can then analyze the information and implement additional security mechanisms to eliminate or reduce the vulnerability.

In another example, the scanner may identify a particular vulnerability in a network service. The benefit presented by that service, however, may outweigh the security risks it brings about. As a result, the organization may continue to allow that service despite its known vulnerability to attack. The attack recognition and response software can detect an attack through that vulnerability and enable the organization to react before the attack compromises the network.

Requirements for effective attack recognition and response

Attack recognition and response software must meet a number of requirements to provide truly effective protection against attacks. The major requirements include:

- *Real-time operation.* The attack recognition and response software must be capable of detecting, reporting, and reacting to suspected attacks in real time. Software that merely logs events and provides audit logs for examination after-the-fact is ineffective. After-the-fact detection is like a burglar alarm that goes off long after the burglar has fled. In addition, many attackers erase logs during the break-in, so their intrusion cannot be detected by merely scanning an event log.

- *Capable of update.* Just as there is a continual launch of new computer viruses, hackers continually find new ways to break into computer systems. As a result, attack

recognition and response software must be capable of continually adding to its knowledge base of known break-in patterns and unauthorized activity.

- *Run on popular network operating systems.* The software must support existing network infrastructures. That means it must support existing network operating systems, such as UNIX and Windows NT.

- *Easy to configure.* Configuration should be easy, without sacrificing effectiveness. The attack recognition and response software should provide a default configuration so that administrators can deploy it quickly and optimize it over time as information accumulates. In addition, the software should provide sample configurations to guide administrators in setting up the system.

- *Easy to manage.* Rapidly rising network management costs present a significant problem for organizations. Attack recognition and response software must be easy to manage so that it does not contribute to this problem. Management of the software over the network from a central location is essential. In addition, the software should be easy to integrate with the existing network management infrastructure. This requires compliance with network management standards such as SNMP.

- *Adaptable to changing security policies.* Today's business environment is dynamic. Organizations are continually changing, driven by many factors, including reorganizations, mergers, and acquisitions. As a result, security policies are also in flux. To remain effective, attack recognition and response software should be easy to adapt to changing security policies. This ensures that these policies can be implemented in fact as well as on paper.

- *Nonobtrusive.* The software should operate in a nonobtrusive way. That is, it should not degrade network performance. It should be transparent to authorized users so that it does not hamper productivity. In addition, it should not alert the intruder to its presence.

The RealSecure solution

RealSecure, from Internet Security Systems, is a real-time monitoring, attack recognition, and response system. It monitors packet flow over a network in real time and analyzes packets for known attack patterns and unauthorized activity using a rule-based approach. When it detects an attack, it reacts automatically according to its configuration. RealSecure can react in four ways when it detects an attack:

- It can alert appropriate personnel through administrator's console messages, e-mail, or pager alerts.
- It can log the event along with associated information.

- It can kill the event by terminating its connection.
- It can initiate a user-supplied script.

Organizations can strengthen security with RealSecure in three primary ways:

- They can use it as an effective, second line of defense behind firewalls, intercepting and disposing of attacks that get through the firewalls, or that originate inside the firewalls.
- They can use it as a means of measuring the effectiveness of the current network security mechanisms by testing whether they are indeed keeping out what they are supposed to be keeping out. For example, RealSecure can detect when a Telnet session is being established from the Internet although the firewall has been configured to disallow Telnet sessions from the Internet.
- They can use it in conjunction with a security assessment package to provide feedback on risks they have accepted in order to provide access to a service that users need but that makes the network vulnerable to attack. RealSecure can determine the number of attempted break-ins through that vulnerability. If the number of attempted break-ins is high, the organization may change the decision to make that service available.

RealSecure provides useful reports on its findings to help organizations assess and tighten their security.

Advanced architecture

RealSecure consists of three components:

- *Recognition engine*. This component monitors the network in real time, detecting and reporting attacks. It reports events to the Administrator's Module.
- *Response Engine*. This component reacts automatically to recognized attack events, triggering prespecified actions ranging from logging attacks and alerting the administrator to terminating offending connections.
- *Administrator's Module*. This component provides GUI (graphical user interface) management of the Recognition and Response engines. The Administrator's Module can monitor and manage all recognition and response engines from a single GUI, simplifying network management.

Recognition engine

The Recognition engine leverages Internet Security System's in-depth knowledge of how systems are attacked. The Engine can detect low-level IP attacks, such as those using IP spoofing, IP fragmentation, or SYN flooding. These low-level forms of attack can bypass packet filter firewalls. The Engine can also detect high-level attacks, such as from Web, FTP, NFS, NIS, or e-mail sessions.

The administrator can configure the Recognition engine to implement specific security policies. The engine can react to (or ignore) connections based on specific packet types, source, and destination IP addresses or address ranges, port numbers, or particular types

of attack patterns. This flexibility enables the administrator to set up custom monitoring for individual hosts and networks. Because the Recognition engine is a passive monitor, operating like a sniffer with built-in security knowledge, it doesn't degrade network performance.

As Figure 1 shows, an organization can place multiple Recognition engines in strategic locations in its network topology. An organization can also use multiple Recognition engines in parallel at a single location to accommodate high bandwidth connections, such as T3 access to the Internet.

Figure 1
The RealSecure real-time attack recognition and response system

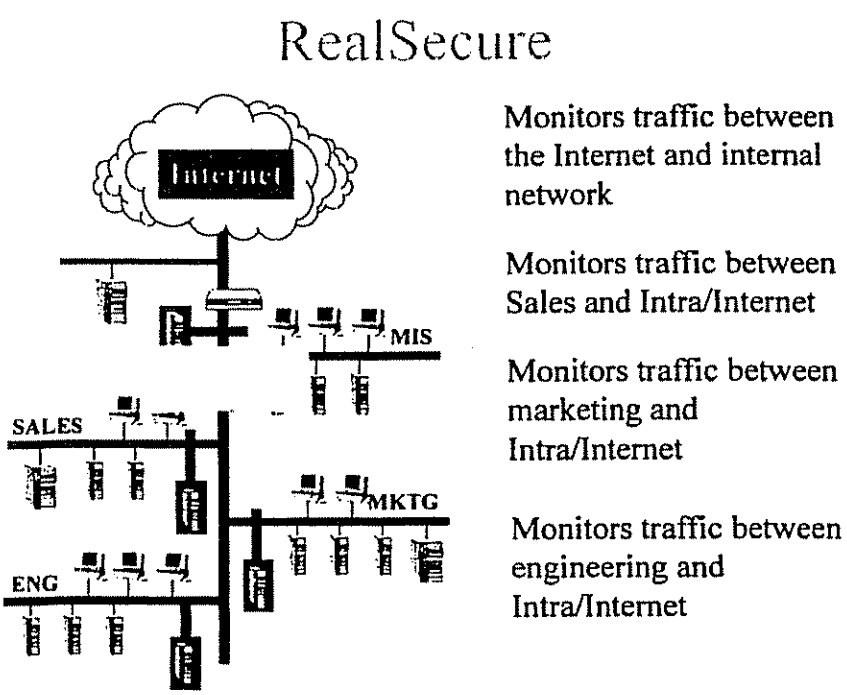

Response Engine

The Response Engine is flexible and can be configured to react automatically to detected attacks in a variety of ways:

- *Alert appropriate people.* The Response Engine can alert the administrator through the Administrator Module, it can send e-mail messages, and it can activate pagers.

- *Log attack.* The Response Engine can be configured to log the attack. It can log the event only or the entire attack session, including all the hacker's keystrokes. The administrator can play back the session later, in its entirety, through an easy-to-use, VCR-like GUI control panel.
- *Terminate session.* The Response Engine can be configured to reset the connection on both the attacker's machine and the target machine when it detects an attack. It terminates the session by sending a reset (RST) packet to the attacker's machine, and it sends an RST packet to the target machine, spoofing the attacker's IP address.
- *Initiate user-supplied scripts.* The response can also be customized with user-supplied scripts that are activated when an attack is detected. In this way, reaction can be tailored to the specific needs of the organization.

The Response Engine's ability to react automatically provides proactive security, without requiring administrator intervention.

Administrator's Module

The Administrator's Module provides a single point of management and control for all Recognition and Response engines in the network. The administrator can configure the Recognition and Response engines from the Administrator's Module. In addition, the Module collects and presents events reported by the Recognition engines in an easy-to-read GUI display.

When a Recognition engine detects an attack, it reports it to the Administrator's Module. The module displays the event in real time, as it is happening, optionally including the hacker's keystrokes and a copy of the hacker's screen. To allow easy, attack events are classified and displayed by high, medium, or low priority. (See Figure 2.)

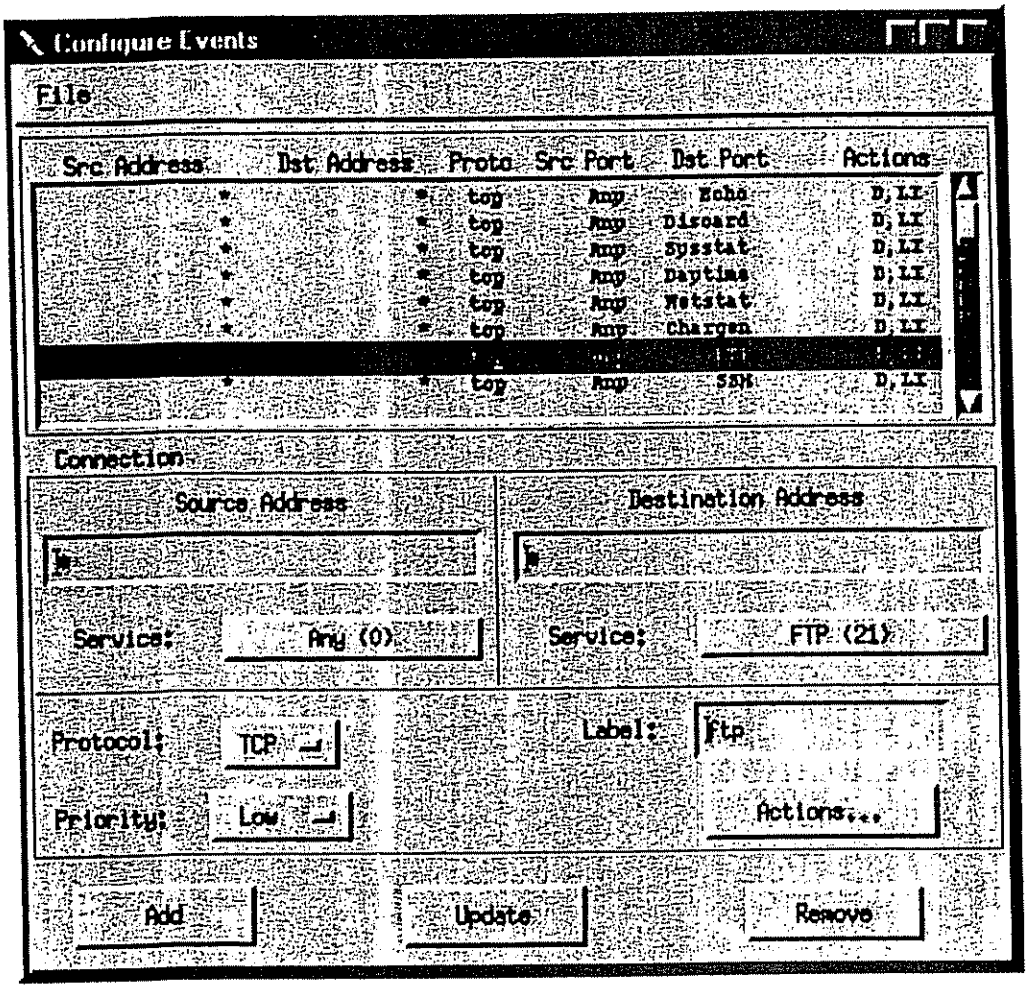

Figure 2
The RealSecure Administrator's Module

Manual response

In addition to automatic response, the administrator can respond manually to reported attacks in a variety of ways using the Administrator's Module GUI:

- *Request additional information.* The administrator can request the Engine to provide more detailed information on the reported attack. This information can include the packet source and packet data, such as e-mail headers.

- *Instruct RealSecure to log the event.* If automatic event logging has not been configured for this event, the administrator can request RealSecure to log the reported event.
- *Instruct RealSecure to kill the event.* If automatic session termination has not been configured for this event, the administrator can request RealSecure to terminate the session.

The administrator can combine automatic and manual response to maintain the exact level of control required.

Easy configuration

The administrator configures RealSecure through the easy-to-use Administrator Module GUI. The configuration specifies the types of checks to be performed by the Recognition Engine and the response to be initiated by the response Engine for each type of attack detected. Through the GUI, the administrator can custom-tailor a security model of the network to match the organization's security policy.

RealSecure includes a default configuration to allow an organization to get up and running quickly. The default configuration is biased towards tight security to ensure the protection of the monitored network. RealSecure also includes a variety of sample configurations at different levels to provide starting configurations for various types of security policies.

Meaningful reports

RealSecure can generate meaningful reports from its event log files. These reports can include such information as the amount of data processed by a Web server each day, or the number of connections that were killed each day and from whom. The Administrator Module can display these reports in graphical form, such as bar or pie charts, for easy review and analysis. (See Figure 3.)

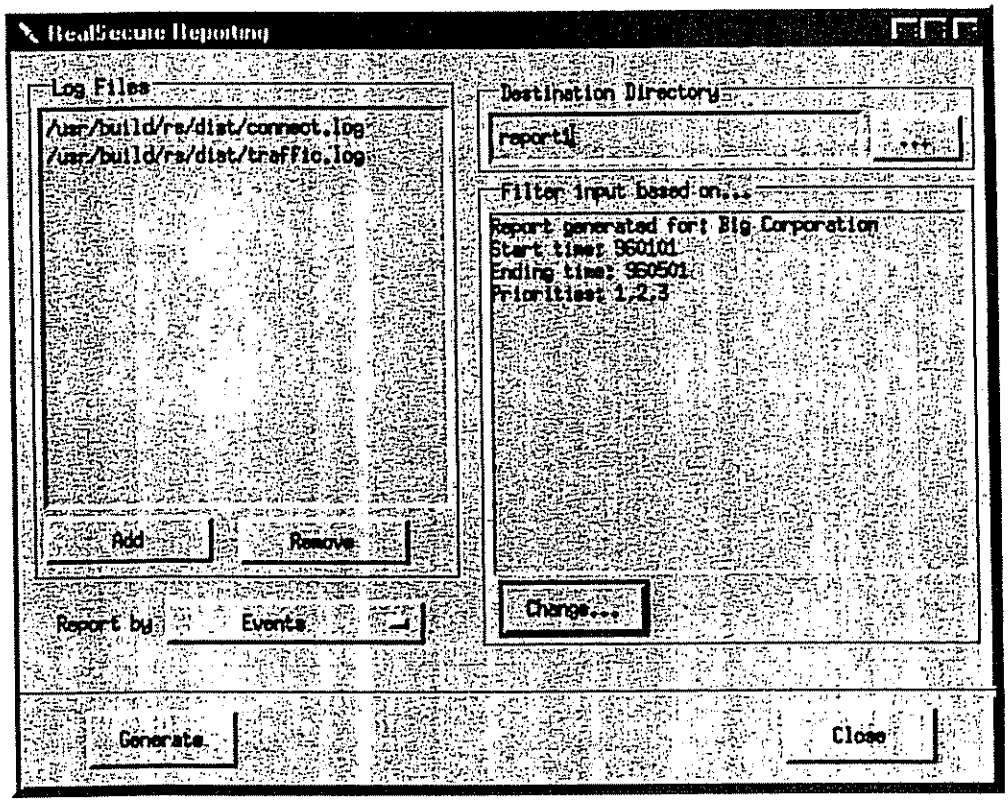

Figure 3
Typical RealSecure Report

Administrators can use these reports to optimize security. The suggested method of implementation is to start with overly tight filtering using the default configuration. The administrator can use the information in the reports to tune the filtering over time as he or she gains better understanding of normal network activity. This iterative tuning reduces the number of alerts without relaxing security.

Conclusions

In order for organizations to compete effectively in today's business environment it is essential that they increase their use of networks, including:

- Expanding the reach of their internal networks by interconnecting LANs into WANs,.
- Opening their internal networks to outside organizations to gain higher levels of interaction with their business partners.

- Taking full advantage of the power of the Internet.

However, the increasing use of networks brings with it increased vulnerability to network break-ins. Traditional network security solutions such as firewalls, operating system security mechanisms, and encryption protect internal networks to a large degree, but hackers still manage to penetrate. Break-ins—whether internal or external—can be costly, and expose sensitive and confidential information to the outside world.

RealSecure augments existing security mechanisms and helps reduce the gap between an organization's security policy and its actual security practice. It enables organizations to measure the effectiveness of their current network security mechanisms in implementing security policy. It also provides an effective second line of defense behind existing mechanisms.

With the combination of RealSecure and traditional security mechanisms, organizations can continue to expand their use of networks, without increasing their risk of network break-ins. In this way, they can maintain their competitive edge while also keeping their security practices in line with their security policies.

About Internet Security Systems, Inc.

Internet Security Systems, Inc. (ISS) is the leading supplier of network security assessment tools, providing comprehensive and innovative audit, monitoring, and response software. The Atlanta-based company's flagship product, Internet Scanner, is the leading commercial attack simulation and security audit tool used to facilitate continuous network security improvement in corporations, financial institutions, and government agencies worldwide. The ISS SAFEsuite family of products provides a comprehensive security framework specifically designed to assess a variety of network security issues confronting web sites, firewalls, servers and workstations. For more information about ISS and its products, contact the company at (770) 395-0150 or visit the ISS web site at http://www.iss.net.

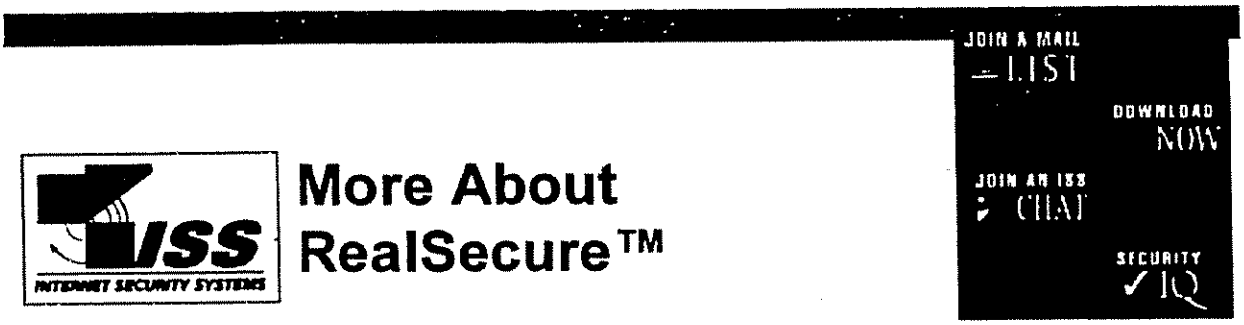

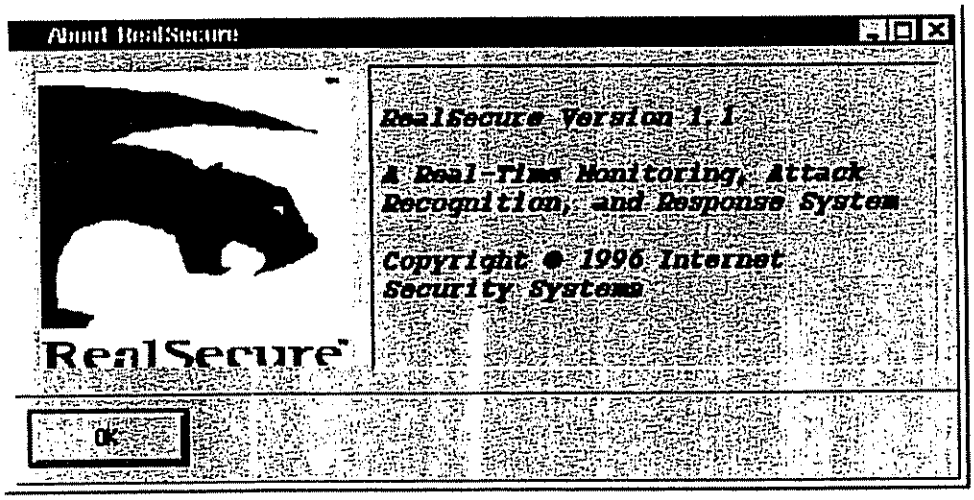

General Description and Comparison to Existing Systems

RealSecure™ was designed to manage the large amounts of data processed by a network in a manner not normally utilized by other network security products. Its configurability and many features make it useful for everything from logging hacking attempts to providing a second line of defense behind a firewall. It can detect and then report network events through various means including email or a user-written program. It can actively defend against attacks by closing attempted connections to your network. What really makes it different from every other network monitoring product is that it is designed to be an enterprise-wide solution, monitoring the company's network at many different points of failure 24 hours a day, 7 days a week.

RealSecure™ is composed of two parts: a filtering engine that watches and actively manages the network and a GUI front-end that reports events and allows the user to configure the engine's scope. Multiple engines can be run on machines near critical points in the network, such as the firewall or a sensitive LAN. The engines report interesting data back to the GUI and automatically handle certain

Linux (right now) and doesn't require too much in the form of system resources. Lots of disk space is good for allowing logging, and more performance is better if the network traffic to be monitored is heavy. The GUI can run on the same platforms, but requires more performance and a graphical display. This makes the GUI best suited for use on a dedicated administrative machine. In fact, it is recommended that any machine running the GUI or the engine be dedicated to that purpose for performance and security reasons.

Starting the GUI shows a short intro screen and then presents the user with the startup window. This window is used to start engines, configure them, and retrieve log files from engines. To start an engine, the user simply types in the name or IP of the machine on which to run the engine. By clicking on an engine in the list, he can also change the configuration info for that engine, shut it down, or fetch files from that machine.

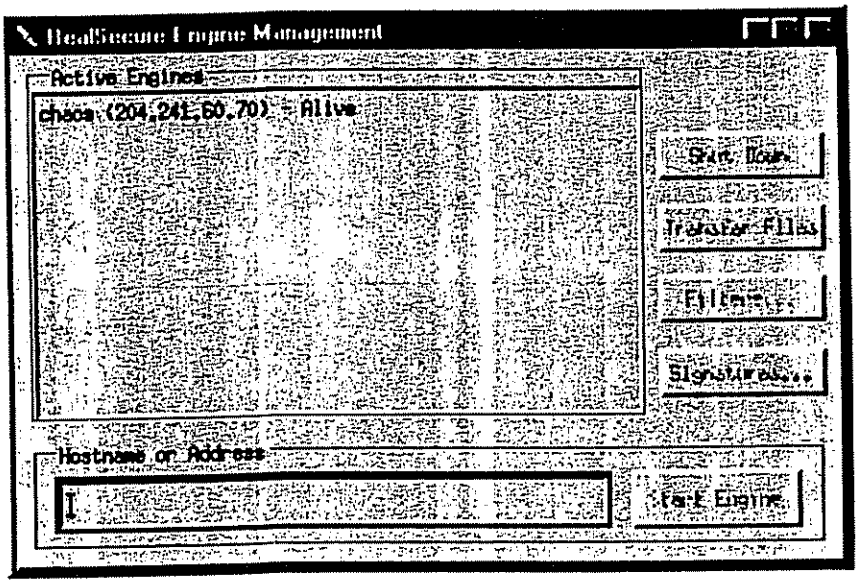

Figure 1 - RealSecure™ Configure Engines Screen

When the user starts an engine, he is presented with a screen which shows a short summary of events classified by security priority: High, Medium, and Low. At the bottom of the screen is a bar graph showing the packets received per second, with the label indicating the maximum ever received. The top menu bar has buttons to load a config file, pop up the GUI configuration editor, or get help.

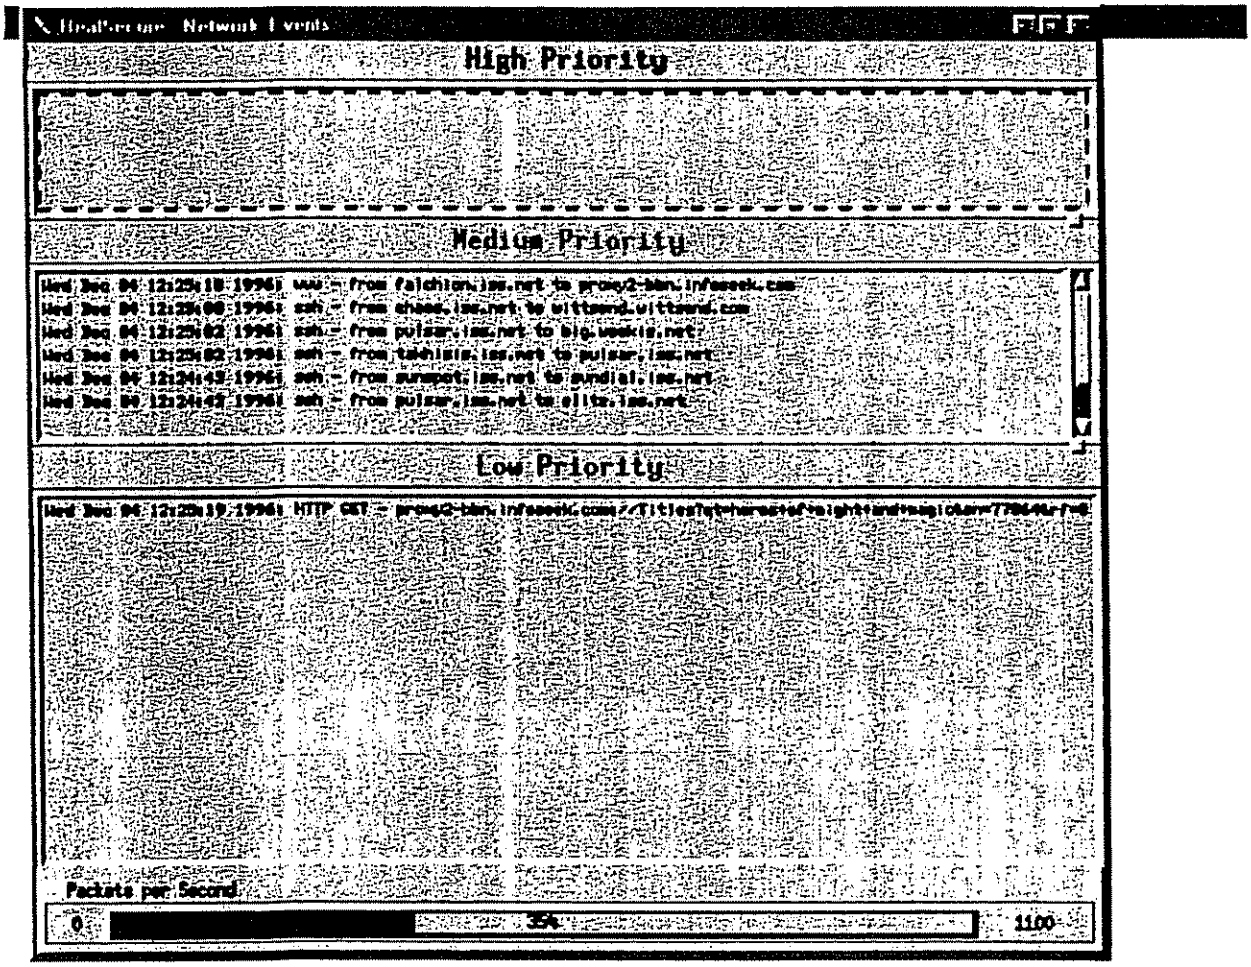

Figure 2 - RealSecure™ Main Screen

As events occur on the network, the engines send messages to the gui to indicate the event and the level. The gui displays them according to priority. The user can double-click on an event in any of the three windows, and then select some action to be performed on that event. At the moment, he can select "More Info", "Log", "View", or "Kill". What that action actually does is limited to the type of event. For instance, if a user decided to log all name queries (DNS), and then tried to kill one of them, it wouldn't be possible because the query would already be finished and there would be no connection to close.

Bringing up the "More Info" window queries the engine for more information about the selected event. This can include items like the source of the packet or perhaps some of the data from the packet like email headers or other important data to log. Its use is to allow the administrator to make a decision whether to ignore the event or take some kind of action.

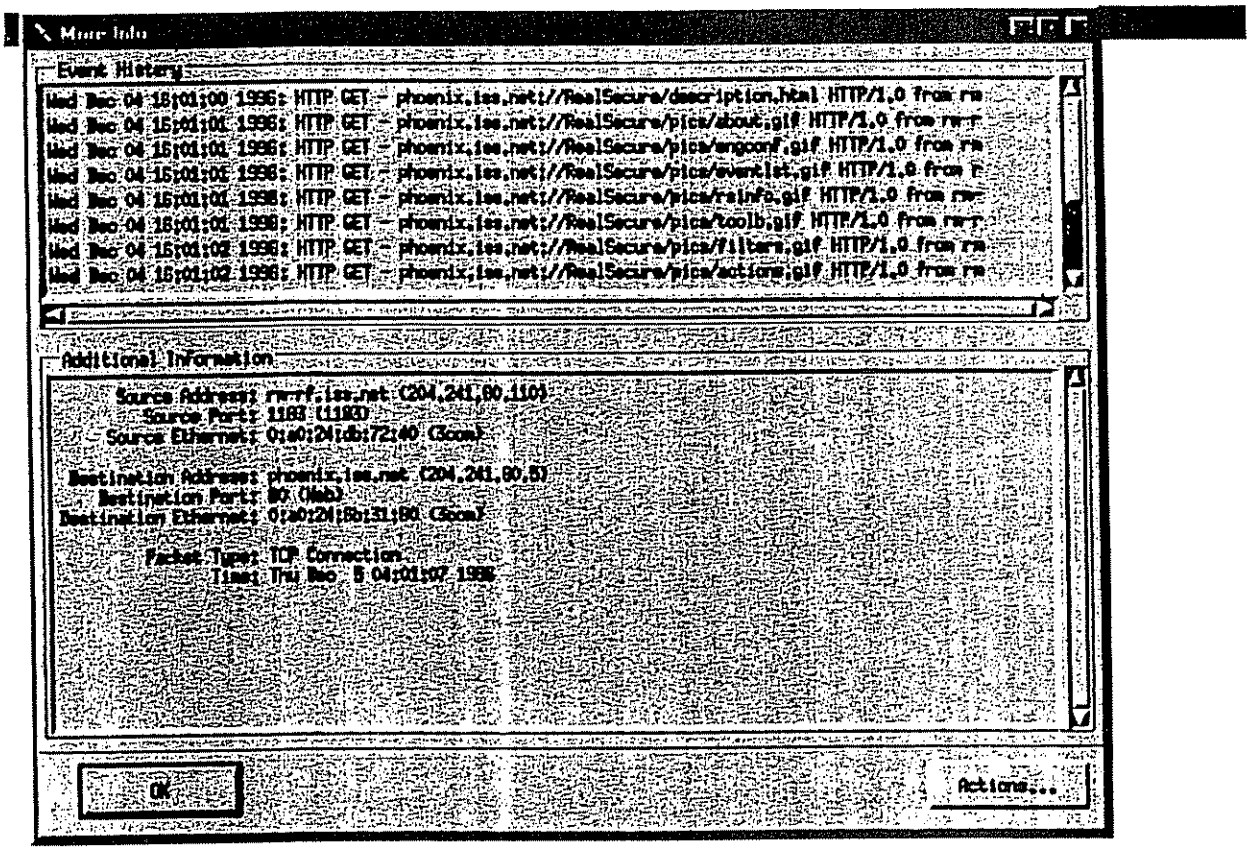

Figure 3 - More Info Screen

The "View" window decodes the data for a given connection into a nice, usable format appropriate for the given connection. For starters, it will show a telnet, rlogin or interactive rsh session in a window with the user's keystrokes in the bottom window. What's nice about this is that the screen looks exactly like what the hacker sees, making it simple to make a decision whether or not to take further action. This can also be used to help network users when they are having problems. The administrator can go on to log or kill the session from this window. The view window looks the same as the playback window described below.

Logging data can be done automatically or manually triggered by the user. Data can be saved in raw format for later playback, or just the pertinent data like where the connection is from can be logged. Logs can be in text format as well for easy perusal or insertion in reports.

The toolbar appears with the engine config screen. It allows quick access to the main components of RealSecure.

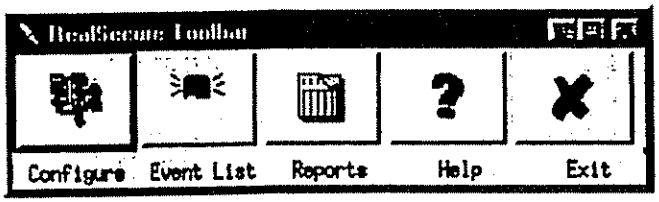

The most important feature of RealSecure™ is its configuration screen, which allows the user to easily tailor a model of the network and prepare custom checks and responses.

```
# Format:
#     Source Address    Dest Address      S Port  D Port Tag    Pri    Action
tcp         129.65.0.0/16     10.0.1.8/32         0       23     Login    1     K,LR=data
```

Figure 5 - Sample filter configuration entry

While this may seem complex, it's easy to understand with a little explanation. This rule takes all TCP packets from 129.65.* (any port) sent to the single machine 10.0.1.8 (port 23) and kills the connection and logs all the data it can about the connection. It is assigned a priority of "high", meaning its notification will appear in the high priority window of the main screen (see Figure 1). The tag will also appear in the window, allowing an administrator to easily pick out which events fit the same categories.

The source and destination address format may be unfamiliar to those who have never configured a firewall. It has two components: the address to match and the number of bits to match from it. All IP's are 32 bits, with each of the dotted numbers signifying eight of those bits. So, 205 is eight bits, 205.16 is sixteen bits, 205.16.18 is twenty-four bits, and lastly, 205.16.18.2 is thirty-two bits. To match all machines coming from 205.16.18.X (where X is some number), the rule would be 205.16.18.0/24. This rule would match 205.16.18.1 and 205.16.18.233, but not 205.16.17.1.

Remember that most users will not have to deal with this config file format, but it's there for those who want to customize the filters very specifically for their network. Most users will probably be using the GUI config, which allows much of the same functionality, but in a bit more friendly manner. For instance, the rule above would be created in the GUI configuration by selecting your network address as 10.0.1.*, selecting a source address of 129.65.*, saying to log and kill all TCP connects to your telnet port.

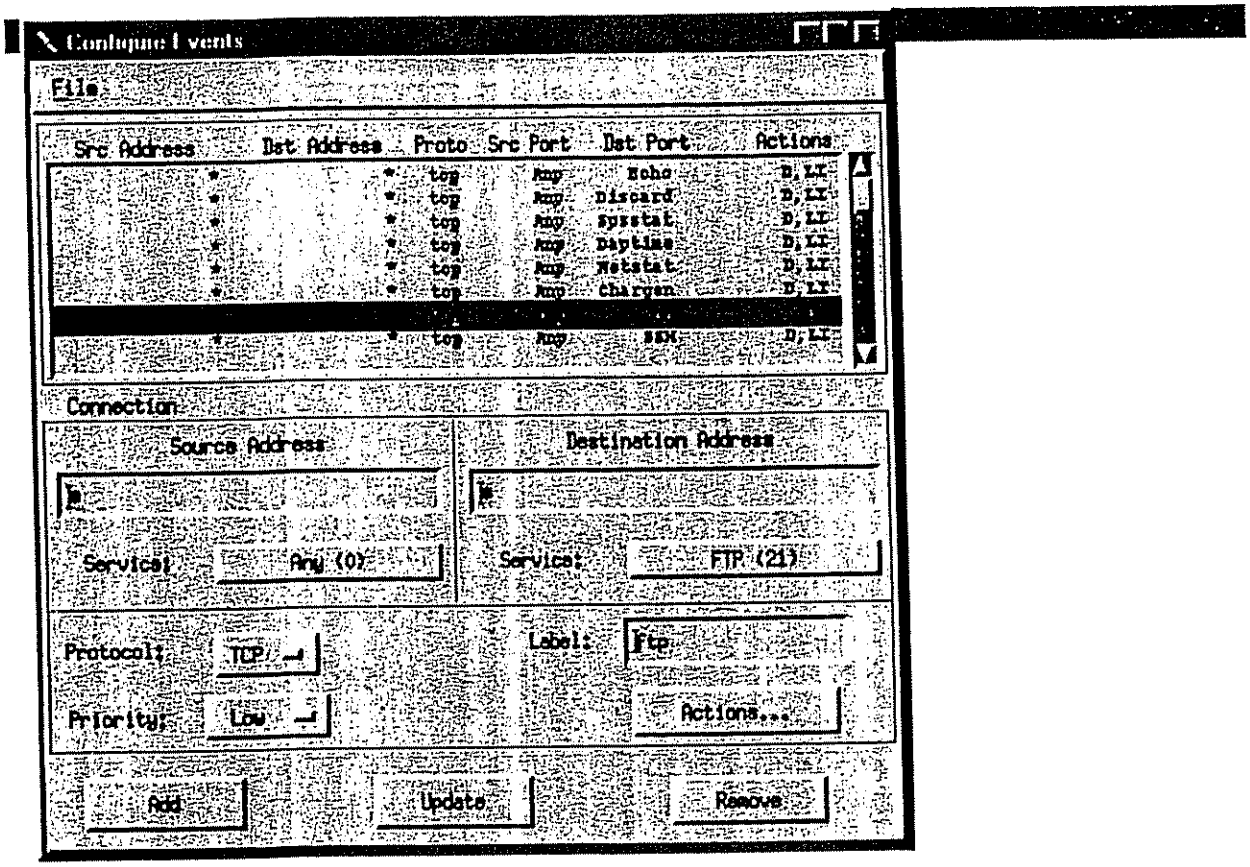

Figure 6 - Filters Configuration Screen

Both the attack signatures config and the filter rules config have a certain set of actions (responses) available for the user to use interactively or the engine to use automatically.

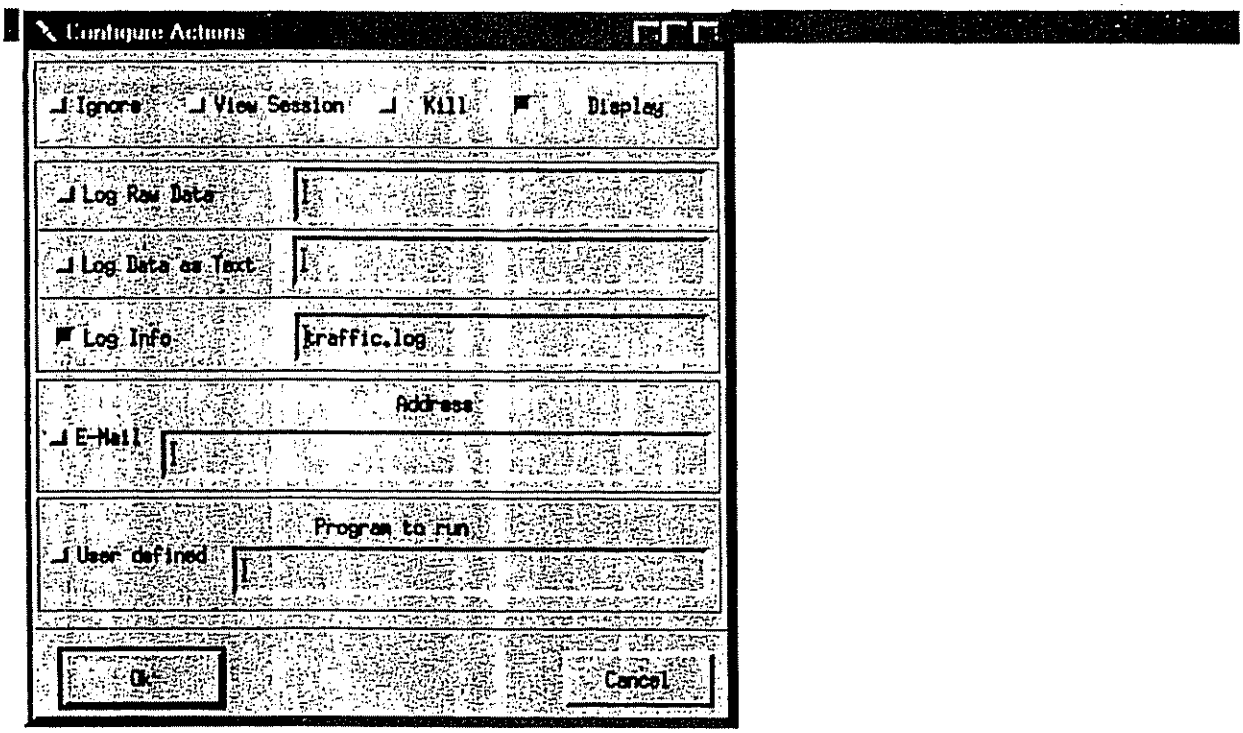

Figure 7 - Actions Configuration Screen

All the different attack signatures are configured through this next screen. Signatures can be enabled or disabled, the log messages can be editted, and any associated actions can be configured. Also, any signature which has tunable parameters (to prevent false alarms) can be managed from this screen.

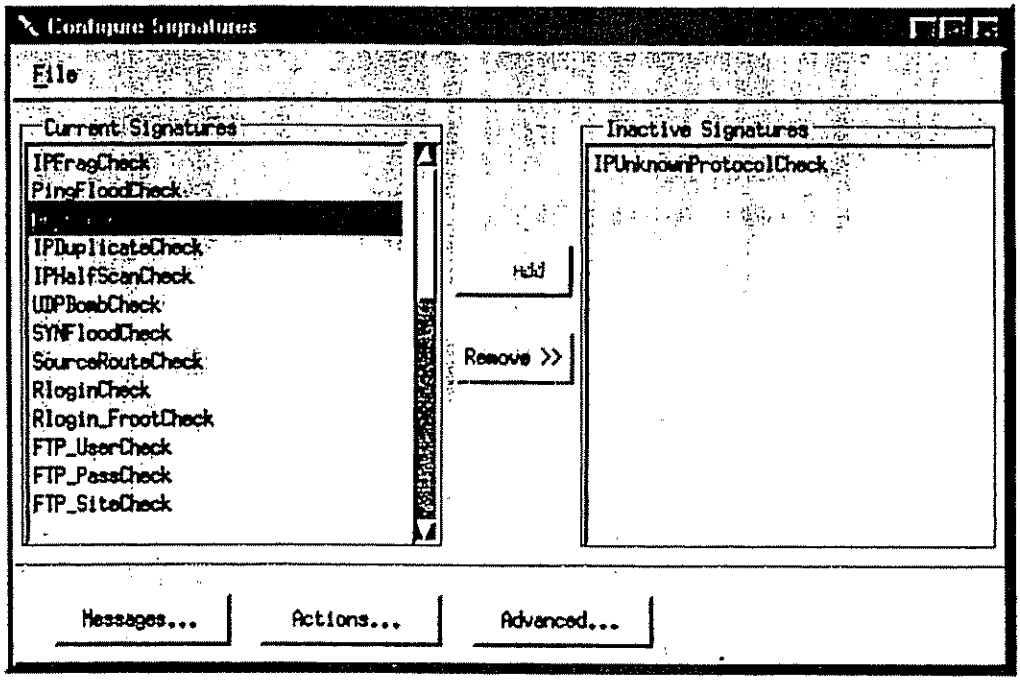

Users can configure the messages and the priority that they appear at by using the config messages dialog.

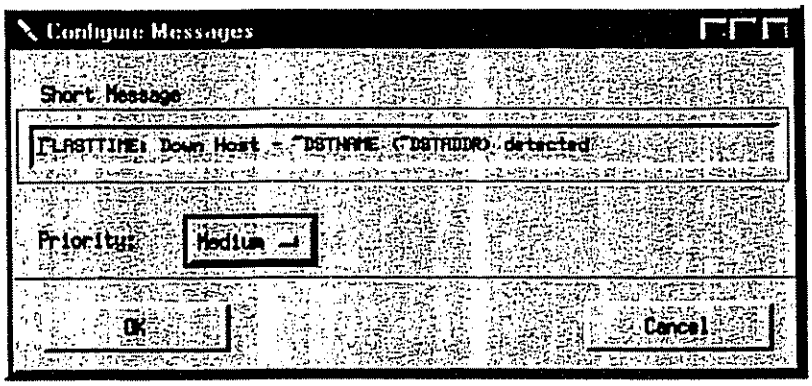

Figure 9 - Messages Configuration Screen

Some signatures have advanced config options which allow them to be custom-tuned for a network. This prevents false positives and makes RealSecure™ a more effective monitoring tool. Here is a picture of a user configuring the SYN Flood Check.

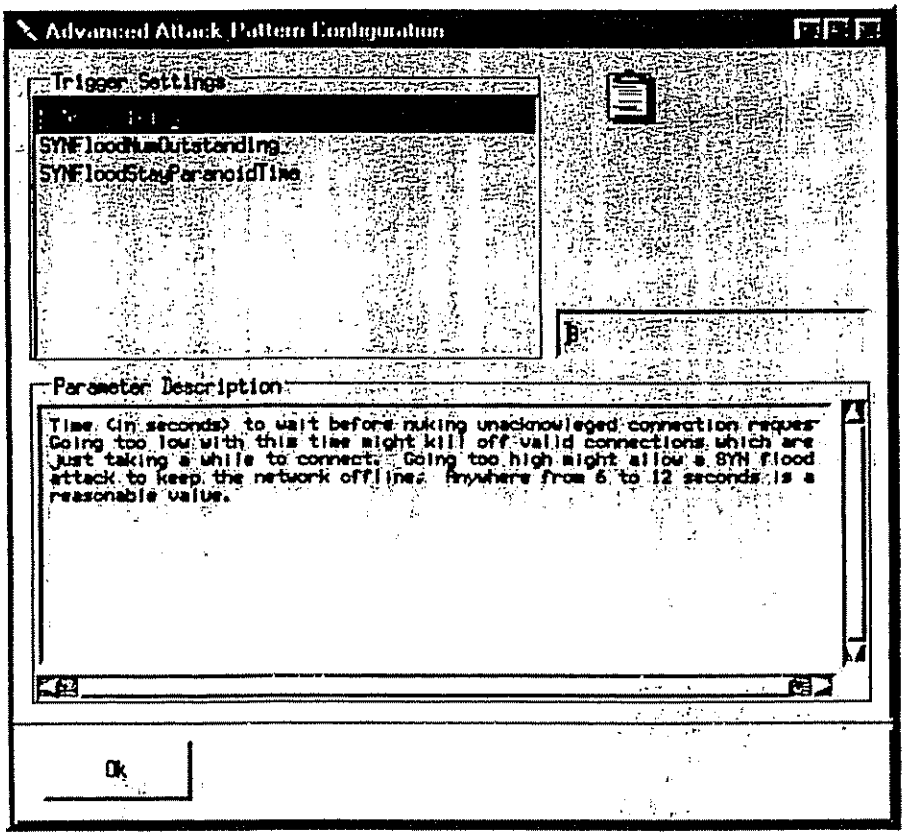

Figure 10 - Advanced Signature Configuration Screen

and then play back the session in real-time. The playback can be paused at any time, or sped up and slowed down. The seek buttons allow the user to jump to the next recorded session. The screen also displays informational messages about the connection.

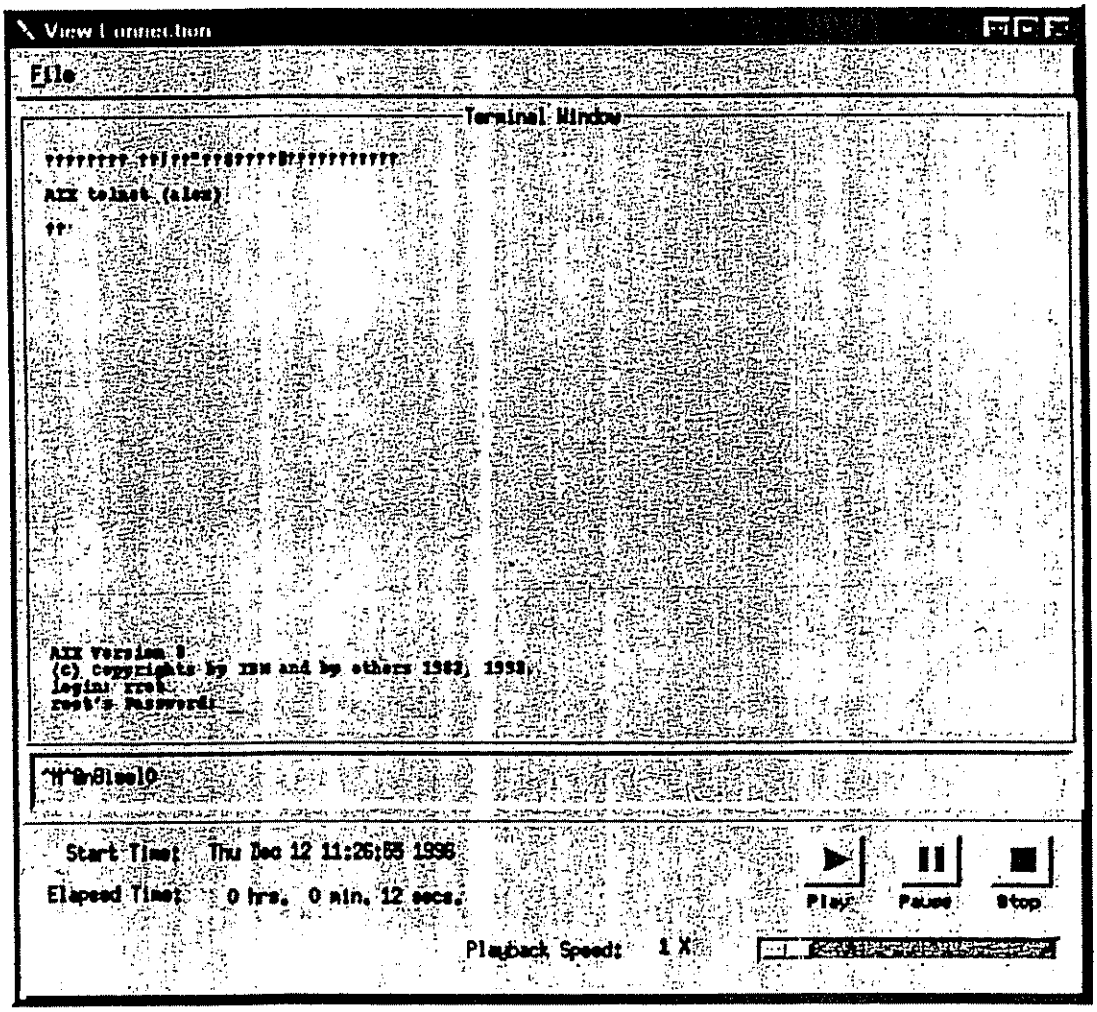

Figure 11 - Playback Screen

The report generation screen allows the user to choose the type of report he wants, as well as what files to use to generate the report.

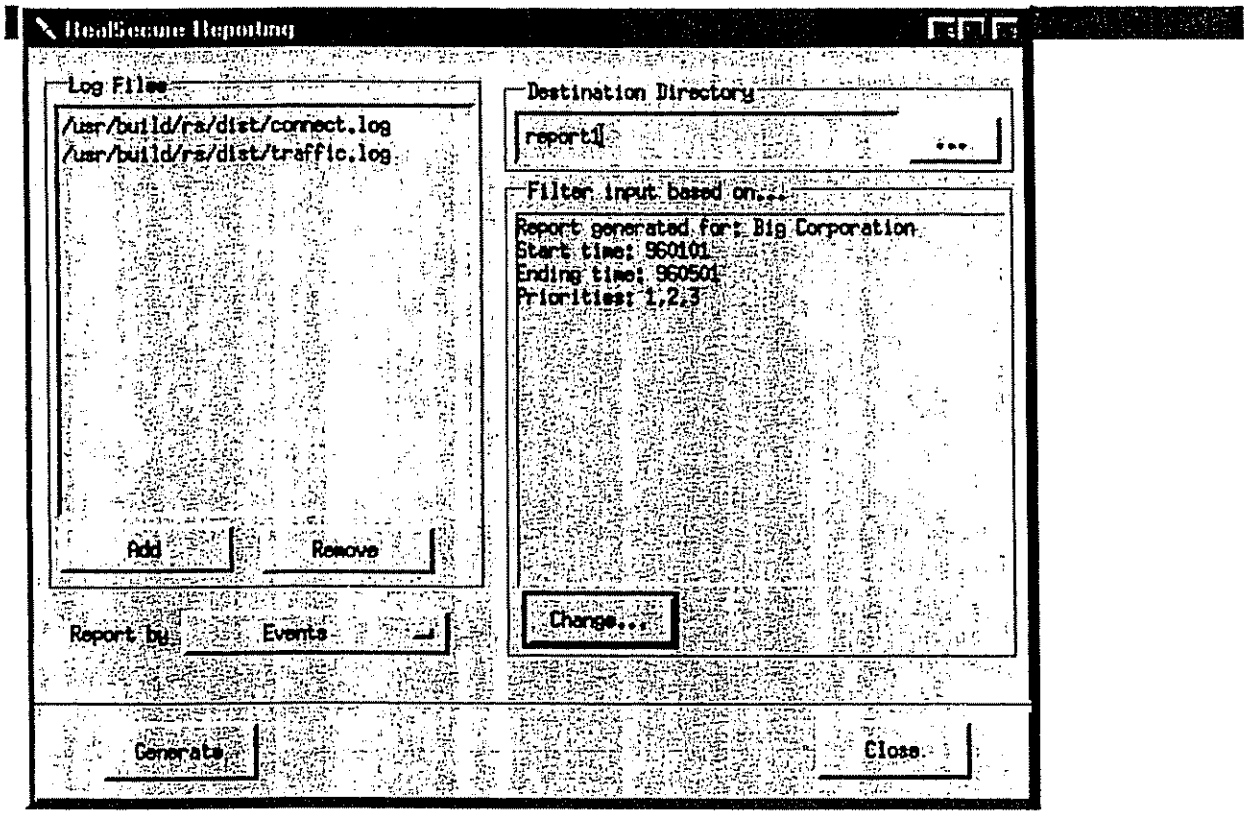

Figure 12 - Report Screen

41 Perimeter Center East · Suite 660 · Atlanta, GA 30346 · (770) 395-0150 · (770) 395-1972 FAX

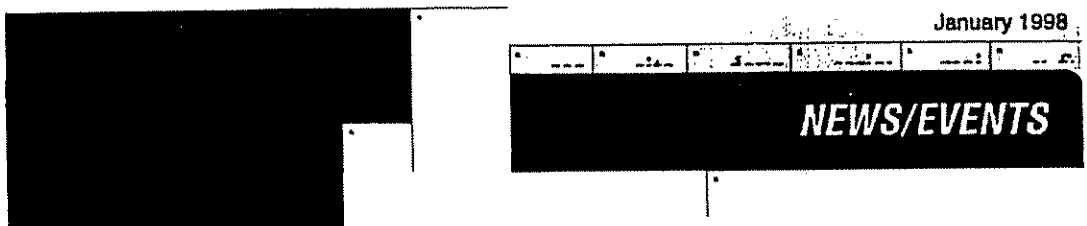

Internet Security Systems Ships RealSecure For Windows NT, Industry's First Real-Time Attack Recognition and Response Tool For Windows NT

-- ISS delivers solution for guarding critical network data around-the-clock and minimizing network security risks --

ATLANTA, Ga., August 19, 1997 -- Internet Security Systems Inc. (ISS), the world's leading provider of network security assessment and monitoring tools, announced today the shipping of RealSecure 1.0 for Windows NT. This unique and powerful network security software helps protect organizations 24 hours a day, 7 days a week from potentially devastating security breaches.

Utilizing ISS' extensive knowledge of security vulnerabilities and its unique patent-pending attack recognition engine, RealSecure detects and intercepts security breaches that penetrate networks and immediately responds to these attacks before a computer and a network are compromised.

"For the first time, Windows NT administrators now have a powerful security tool for protecting their networks from invasion," said Patrick Taylor, director of product management for ISS. "Installing RealSecure at critical locations on a network, such as behind a firewall, enables companies to confidently conduct business online. System administrators can rest easy knowing RealSecure is guarding their network 24 hours a day and responding in real-time to any suspicious network activity."

Much like a security camera in a building lobby, RealSecure provides a fundamental layer of protection to traditional network perimeter defenses, such as firewalls, and ensures that these defenses are working. RealSecure has the ability to monitor and protect all network devices and services such as machines connected to the Internet, intranets and servers, with no effect on network performance.

RealSecure monitors traffic originating on the internal network or passing through firewall tunnels and responds instantly to suspicious activity by logging the session for later review, immediately notifying the administrator or automatically terminating the connection. In

addition, sessions can be played back at any time for further evaluation or criminal evidence. RealSecure is highly configurable and can be customized to meet the requirements of a company's security policy.

An innovative, distributed architecture allows administrators to easily install RealSecure monitoring engines at critical points on their network and manage these engines from a single, centralized RealSecure console, thereby consolidating management of the overall security of an enterprise network. RealSecure's easy-to-use, Windows-based, graphical user interface (GUI), displays information in the format that best meets the needs of network administrators.

RealSecure contains fully-functional reporting facilities that allow administrators to develop meaningful reports. These reports can include such information as the amount of data processed by a Web server each day, or the number of connections that were killed and from whom. The RealSecure management console can display these reports in graphical form, such as bar charts, for easy review and analysis.

RealSecure is a critical component of ISS' SAFEsuite product family. Together, SAFEsuite's comprehensive set of network security assessment and monitoring tools enable organizations to minimize and manage their network security risks.

Pricing

Shipping now, RealSecure carries a U.S. suggested list price of $4,995 for a single perpetual license and can be purchased directly from ISS or through authorized ISS Security Partners worldwide. A UNIX version of RealSecure is also available.

About Internet Security Systems

Internet Security Systems, Inc., (ISS) is the pioneer and world's leading supplier of network security assessment and monitoring tools, providing comprehensive software that enables organizations to proactively manage and minimize their network security risks. ISS' SAFEsuite product family automatically detects, monitors, and responds to the growing number of network security vulnerabilities and threats. The Atlanta-based company's flagship product, Internet Scanner, is the world's leading security auditing tool used to eliminate network security vulnerabilities in corporations, government agencies, and financial institutions including 9 out of the top 10 U.S. banks. ISS' real time attack recognition and response tool, RealSecure, is the leading network monitoring software used to automatically guard networks from external threats and internal misuse. For more information, contact the company at (800) 776-2362 or (770) 395-0150 or visit the ISS Web site at http://www.iss.net.

###

Sales Inquiries: sales@iss.net

41 Perimeter Center East · Atlanta, GA 30346
(770) 395-0150 · (770) 395-1972 FAX